

Efficient Instance Retrieval with Standard and Relational Path Indexing

Alexandre Riazanov and Andrei Voronkov

Department of Computer Science
University of Manchester
Preprint Series
CSPP-20

Efficient Instance Retrieval with Standard and Relational Path Indexing

Alexandre Riazanov and Andrei Voronkov

January 31, 2003

Abstract

Backward demodulation is a simplification technique used in saturation-based theorem proving with superposition and ordered paramodulation. It requires instance retrieval, i.e., search for instances of some term in a typically large set of terms. Path indexing is a family of indexing techniques that can be used to solve this problem efficiently. We propose a number of powerful optimisations to standard path indexing. We also describe a novel framework that combines path indexing with relational joins. The main advantage of the proposed scheme is flexibility, which we illustrate by sketching how to adapt the scheme to instance retrieval modulo commutativity and backward subsumption on multi-literal clauses.

Keywords: backward demodulation, instance retrieval, path indexing.

Copyright © 2000, University of Manchester. All rights reserved. Reproduction (electronically or by other means) of all or part of this work is permitted for educational or research purposes only, on condition that no commercial gain is involved.

Recent preprints issued by the Department of Computer Science, Manchester University, are available on WWW via URL <http://www.cs.man.ac.uk/preprints/index.html> or by ftp from <ftp://ftp.cs.man.ac.uk> in the directory `pub/preprints`.

Contents

1	Introduction	1
2	Preliminaries	2
2.1	Instance Retrieval	2
2.2	Term Indexing	2
2.3	Discrimination Trees	3
2.4	Standard Path Indexing	4
2.4.1	Paths in Terms	4
2.4.2	General Scheme for Path Indexing	5
2.4.3	Implementation	5
2.5	Experiment Settings	8
3	First Experiments	9
4	Better Path-List Representation	10
5	Rearranging the Order of Path-List Intersection	11
6	Compiled Cleanup	12
7	Relational Path Indexing	15
7.1	General Scheme for Retrieval	15
7.2	Implementations	16
7.2.1	Simple Implementation of Retrieval	17
7.2.2	Early Cleanup	17
7.2.3	Combining Joins with Compiled Cleanup	17
7.3	Experiments	17
7.4	Instance Retrieval Modulo Commutativity	18
7.5	Backward Subsumption on Multiliteral Clauses	20
8	More Experiments	21
9	Related and Future Work	23
	Bibliography	24

1 Introduction

Effectiveness of *saturation-based theorem proving* in first-order logic with equality ([Nieuwenhuis and Rubio 2001]) strongly depends on the use of various redundancy detection and simplification techniques. In provers implementing the superposition or ordered paramodulation calculi the most widely used simplification technique is probably *demodulation*, also known as *term*

rewriting. If we have derived a unit equality $s \simeq t$, we can simplify other clauses in the following way. A clause $C[s\theta]$, where θ is some substitution, can be replaced by a “simpler” clause $C[t\theta]$, if $s\theta \succ t\theta$, where \succ is the used simplification ordering on terms.

Typically, demodulation is applied in two modes. *Forward demodulation* rewrites newly derived clauses by unit equalities that are already in the current clause set. This report deals with *backward demodulation*, which uses newly derived positive unit equalities to rewrite some old clauses. The main ingredient of any implementation of backward demodulation is a procedure for finding out which parts of which clauses can be rewritten by a given equality. This task can be usually reduced to solving the *instance retrieval* problem, i.e., finding all or some instances of a query term in a set of terms. Apart from backward demodulation in theorem proving, instance retrieval can also be used as a component of *completion* in *term rewriting* (see, e. g., [Dershowitz and Plaisted 2001]).

Efficient search in large term sets usually requires some kind of *indexing* ([Graf 1996, Sekar, Ramakrishnan and Voronkov 2001]). *Path indexing* introduced in [Stickel 1989, Ramesh, Ramakrishnan and Warren 1990] is an indexing technique suitable for instance retrieval. The contributions of this report can be divided in two parts. First, we propose and assess experimentally a number of powerful optimisations to the standard implementation of path indexing for instance retrieval. Second, we describe a novel technique that combines path indexing with relational joins in order to provide perfect filtering for instance retrieval. The main advantage of the proposed framework is flexibility, which we demonstrate by adapting it to more complex problems, such as instance retrieval modulo commutativity and backward subsumption on multiliteral clauses.

2 Preliminaries

2.1 Instance Retrieval

The work described in this report is aimed at finding efficient solutions to the following decision problem. Given a set of first-order terms \mathcal{I} and a *query term* q , find all such substitutions θ that the *instance* $q\theta$ of q is in \mathcal{I} . This *retrieval* operation is supposed to be performed frequently. In addition, the database of terms \mathcal{I} can be dynamically updated between calls to retrieval by adding or removing some terms. We are interested in cases when \mathcal{I} grows very large and straightforward enumeration-based search becomes hopelessly inefficient.

It is important to emphasise that upon retrieval we are interested in finding both the substitutions θ and pointers to the datastructures representing the found instances $q\theta$ *explicitly*. In the context of backward demodulation, the substitutions are needed, in the first place, for instantiating the right-hand side of the rewrite rule. Additionally, if we allow demodulation by equalities that are not preordered, the ordering constraint associated with an equality has to be checked on the retrieved substitutions.

2.2 Term Indexing

The mainstream approach to implementing search in large term sets efficiently is based on *term indexing* (see [Graf 1996, Sekar et al. 2001]). Devising a term indexing technique for instance retrieval amounts to finding a datastructure for representing the term sets \mathcal{I} , called an *index*. The main purpose of the index is to facilitate fast algorithms for retrieval of instances of various

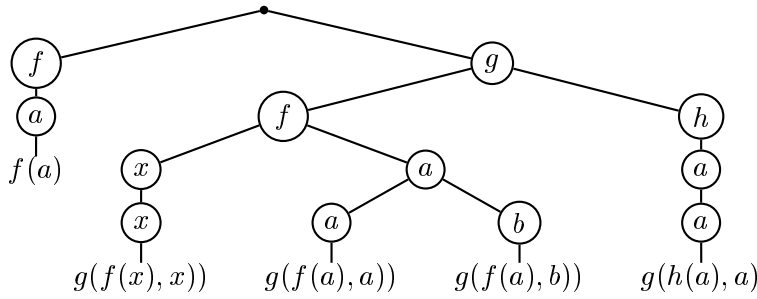


Figure 1: Discrimination tree for $\{f(a), g(f(x), x), g(f(a), a), g(f(a), b), g(h(a), a)\}$

query terms q . Since the database of terms is dynamically changing, we should also provide efficient procedures for *maintenance* of the index, i.e., insertion and deletion of terms.

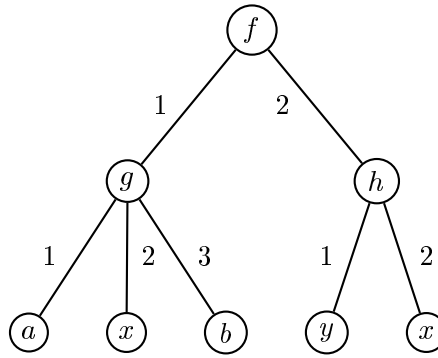
In what follows, we only consider indexing techniques that provide *filters* for the specified retrieval conditions, i.e., they always find (a superset of) all solutions. Moreover, we are only interested in *perfect filtering*, where all the retrieved terms are solutions.

Shortly we will describe two classical examples of indexing techniques for instance retrieval: *discrimination trees* and *path indexing*. The former will be used as a benchmark in our experiments. The latter provides a basis for the main material of this report. In Sections 4, 5, 6 we describe a number of optimisations which allow very efficient implementation of standard path indexing. In Section 7 we describe a novel technique based on path indexing that extends well to other important problems: instance retrieval modulo commutativity and backward subsumption on multiliteral clauses.

2.3 Discrimination Trees

A *discrimination tree* for a set of terms \mathcal{I} can be viewed as a tree whose leaves contain terms from \mathcal{I} and inner nodes are labeled by term symbols, i.e., function symbols and variables. If a branch leads to a leaf with a term t , the corresponding sequence of the inner node labels from the branch is exactly the *string representation* of the term t , i.e., the sequence of symbols obtained by traversing t in the left-to-right depth-first manner. Figure 1 shows a discrimination tree for the set $\{f(a), g(f(x), x), g(f(a), a), g(f(a), b), g(h(a), a)\}$.

Maintenance of a discrimination tree is straightforward regardless of the chosen specific representation. Retrieval of instances is done by simultaneously traversing the query term and traversing the discrimination tree with backtracking. If the symbol in the current position in the query term is a function symbol, a forward move in the tree can only be done towards a node labeled by this function symbol. If we fail to find such a node, we have to backtrack to the most recent backtrack point. If the symbol in the current position in the query is an uninstantiated variable, the current position in the tree becomes a backtrack point, and we skip some part of the leftmost branch starting with this node. Labels from this part of the branch must represent a term which becomes the current instance of the variable. Upon backtracking, we follow the leftmost of the yet unexplored branches. When a variable instantiated by a term s is met in the query, we treat it as if it were s itself. When a leaf is reached, we report its content as a found instance together with the collected substitution, and backtrack with the purpose of finding more instances. For more details on discrimination trees and related algorithms see [Graf 1996].


 Figure 3: Labeled tree for $f(g(a, x, b), h(y, x))$

2.4.2 General Scheme for Path Indexing

Standard path indexing is based on the following observation: every instance t of a term q *contains all paths* from q . This property allows us not to consider terms that do not have some paths from q , as potential instances of q . We can always restrict ourselves by considering only the maximal paths in q . Indeed, a term t contains all paths from q if and only if it contains all maximal paths from q .

Let us formulate a general scheme for path indexing. The index for a set of terms \mathcal{I} should give us a possibility to extract the set $\mathcal{I}_\pi = \{t \in \mathcal{I} \mid \pi \in \text{Paths}(t)\}$ for any given path π . When we need to retrieve all instances of a term q from \mathcal{I} , we compute the *candidate set* containing all terms from \mathcal{I} that contain all maximal paths from q :

$$\text{CandSet}(\mathcal{I}, q) = \bigcap_{\pi \in \text{MaxPaths}(q)} \mathcal{I}_\pi.$$

The candidate set $\text{CandSet}(\mathcal{I}, q)$ contains all instances of q in \mathcal{I} , and typically is much smaller than \mathcal{I} . Unfortunately, it may contain some terms that are not instances, since the presence of all paths from q in a term t does not guarantee that t is an instance of q . For example, $t = f(g(a, a), h(b, b))$ contains all paths from $q = f(g(a, x), h(y, x))$, but is not an instance of q , since t has different terms a and b in the positions corresponding to occurrences of the same variable x in q . Even if we know that $\text{CandSet}(\mathcal{I}, q)$ coincides with the set of instances, as in the case of a linear term q , it tells us nothing about the corresponding substitution, which we have to find explicitly. In order to resolve these problems, we need to perform additional *cleanup* on $\text{CandSet}(\mathcal{I}, q)$, getting rid of terms that are not instances of q and computing substitutions for those that are.

2.4.3 Implementation

Our implementation of the standard path indexing technique is a straightforward specialisation (for instance retrieval) of the implementation described in [McCune 1992].

The sets \mathcal{I}_π are represented as sorted lists of pointers to the datastructures representing the indexed⁴ terms, or any other objects that can serve as *term identifiers*. These lists will be called *path-lists*. We assume that all indexed terms and their subterms are *perfectly shared*, i.e., different pointers always represent syntactically different terms. Intersection of sorted lists can be computed according to the procedure in Figure 4.

```

procedure intersect( $L_1, \dots, L_n$ )
  assume  $n > 1$ ,  $L_1, \dots, L_n$  are in ascending order
  if  $L_1 = nil$  then stop
   $cand := hd(L_1)$ 
   $i := 1$ 
  repeat indefinitely
    while  $L_i \neq nil$  and  $hd(L_i) < cand$ 
       $L_i := tl(L_i)$ 
    if  $L_i = nil$  then stop
    if  $hd(L_i) = cand$ 
      then
         $L_i := tl(L_i)$ 
        if  $i = n$ 
          then
            report  $cand$  as a member of the intersection
            if  $L_1 = nil$  then stop
             $cand := hd(L_1)$ 
             $i := 1$ 
          else
             $i := i + 1$ 
        else
           $cand := hd(L_i)$ 
           $i := 1$ 
    end
  end

```

Figure 4: Intersection of sorted lists

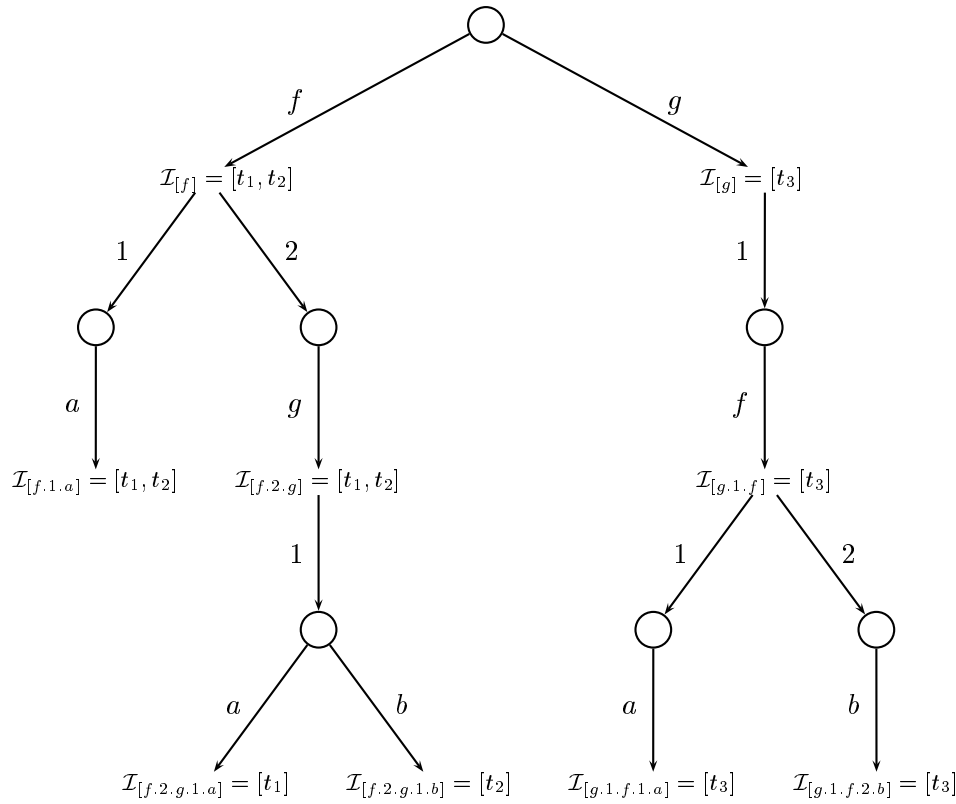


Figure 5: Example of a trie for accessing path lists.

$$t_1 = f(a, g(a)), \quad t_2 = f(a, g(b)), \quad t_3 = g(f(a, b))$$

For accessing the path-list, associated with a path, any hashing in which the path is the key will do. Another solution, which is probably optimal, is to employ the well-known datastructure called *trie*. In our settings, a trie is a directed tree whose edges are labeled by function symbols or non-negative integers. Every node with an incoming edge labeled by a function symbol corresponds to a path π in some terms from \mathcal{I} , and contains the pointer to the path-list \mathcal{I}_π . The correspondence is straightforward: π is the sequence of node labels in the path leading from the root of the trie to the node in question. Consider an example with a small indexed set $\mathcal{I} = \{t_1, t_2, t_3\}$, where $t_1 = f(a, g(a))$, $t_2 = f(a, g(b))$ and $t_3 = g(f(a, b))$. The corresponding trie is depicted in Figure 5.

Given a path π , to find \mathcal{I}_π , we simply navigate through the trie, following the nodes, labeled with the corresponding symbols from π , until we get to the node containing \mathcal{I}_π . If we fail to find such a node in the trie, we assume that \mathcal{I}_π is empty. Maintaining the index is also straightforward. To insert a term t , for every path π in it, we find the node with \mathcal{I}_π and add t to this path-list. If there is no such node in the trie, we first add the corresponding branch to the trie with empty path-lists in the new nodes. To remove t , we delete it from all \mathcal{I}_π , where $\pi \in \text{Paths}(t)$. As soon as a path-list \mathcal{I}_π becomes empty, it is safe to remove the corresponding node from the trie.

Finally, let us describe how the cleanup on the candidate set is implemented. As soon as retrieval procedure for a query q reports a newly found candidate term t , we apply a matching procedure $\text{MatchCand}(q, t)$ to it, in order to check if t is an instance of q and to find the corre-

```

procedure MatchCand( $q, t : \text{term}$ )
   $\theta : \text{substitution} := \emptyset$ 
  if MatchCand'( $q, t, \theta$ )
    then report success with substitution  $\theta$ 
    else report failure
end

fun MatchCand'( $q, t : \text{term}, \text{var } \theta : \text{substitution}$ ) : bool
  if  $q$  is a variable  $x$ 
    then
      if  $x$  is instantiated in  $\theta$ 
        then return ( $x\theta = t$ )
        else
          add assignment  $x \rightarrow t$  to  $\theta$ 
          return true
      else
        let  $q = f(q_1, \dots, q_n), t = f(t_1, \dots, t_n)$ 
        return  $\bigwedge_{i=1}^n \text{MatchCand}'(q_i, t_i, \theta)$ 
  end

```

Figure 6: Matching test for candidate cleanup

sponding substitution. Although *MatchCand* can implement the standard matching algorithm, for the sake of efficiency, in the algorithm shown in Figure 6 we assume that t has all paths from q .

2.5 Experiment Settings

In the rest of the report we compare experimentally several implementations of instance retrieval. We want to perform extensive and reproducible experiments on realistic benchmarks. To fulfill these requirements, we have chosen to use the COMPIT methodology (see [Nieuwenhuis, Hillenbrand, Riazanov and Voronkov 2001]). To make the benchmarks realistic, a real prover is run on real problems, and all operations related to instance retrieval are logged into a special file. To test a particular implementation, it is linked to a simple driver code. The driver reads the benchmark file and submits the index update and retrieval operations, together with all the necessary data, to the implementation in question. It also measures the time, taken by retrieval and index updates, memory, used by the index, collects various statistics on the operations, and prints it all in a form suitable for easy maintaining and processing of large databases of experimental results.

For our experiments, we use all the 2896 equality problems from the TPTP library ([Suttner and Sutcliffe 1997]), version 2.4.1. The benchmarks are generated by running our system Vampire ([Riazanov and Voronkov 2002]), v5.0, on a Pentium III machine, 1GHz, 256Kb cache, running Linux, giving it 200 seconds time limit and 256Mb memory limit. Most of the time, we deal with the 585 pure equality (PEQ) problems from this set, although in Section 8 we give a summary of results on the general class of problems with equality (GEQ). Figure 7 presents the main quantitative characteristics of both benchmark suites. The characteristics are first

	GEQ	PEQ
Total number of problems	2896	585
Average number of insertions	200866	59948
Average number of deletions	162489	22419
Average maximal index size	51897	40964
Average indexed term size	9	12
Average number of queries	15018	33577
Average query size	10	14
Average queried index size	26578	22858
Average number of instances per query	11	25

Figure 7: Benchmark suite profiles

calculated for every particular problem, and the corresponding figures in the table are obtained by dividing their sum by the number of participating problems.

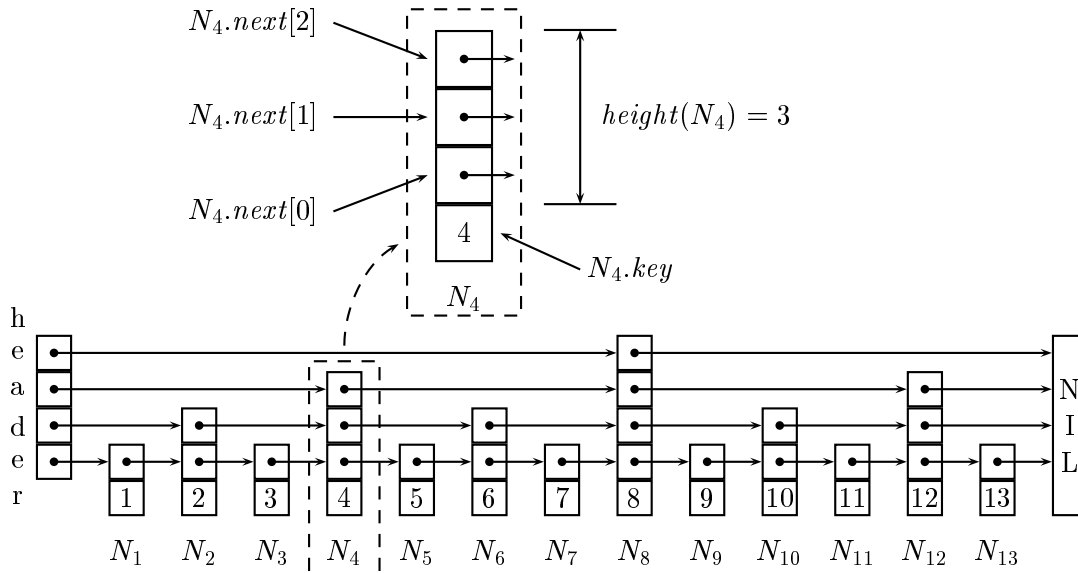
The main parameter of comparison is the overall time taken by both maintenance and retrieval. When two techniques are compared, we compare the sums of their running times on all problems from the benchmark suite in question, thus obtaining an average figure. We also compare memory usage, which is represented by two estimates: the lower and upper bound. The lower bound is the maximal amount of memory occupied by the index datastructures. This figure does not include memory that was used by the datastructures but later reclaimed. This corresponds to an ideal situation when all reclaimed memory can be reused by the index itself or other datastructures in the prover. Usually, this is not the case, and to balance the comparison we also calculate the upper bound. In this case we assume that the reclaimed memory can only be reused by the index itself, and the pieces of memory, reclaimed but not reused, are included in the estimate. Since in all our experiments the margin between the estimates turns out to be negligible (less than 1%), in what follows we only consider the upper bound.

All our implementations of indexing techniques are written in C++. The driver code is written in C, so it can be used for testing implementations both in C and C++. The driver code, some of the implementations of indexing techniques discussed here, and the Vampire source, modified for benchmark collection, will soon be publically available as parts of the COMPIT distribution package [Nieuwenhuis et al. 2001].

3 First Experiments

We start our investigation by comparing the performance of our implementation of discrimination trees (Section 2.3) with the standard implementation of path indexing (Section 2.4.3). On our main benchmark suite, the PEQ problems, the path indexing implementation is inferior to discrimination trees, as it is **3.23 times slower**, while uses 5% more memory.¹ Nevertheless, the idea behind path indexing seems promising, and in the following three sections we describe optimisations that altogether make the standard path indexing technique much more attractive.

¹Here and below, when we discuss experimental data, for details the reader is referred to Section 8, Figures 18 and 19.

Figure 8: Skip list for $[1, 2, \dots, 12, 13]$

4 Better Path-List Representation

Our analysis of several randomly chosen execution profiles for the standard path indexing implementation indicated that, when the index is sufficiently large, most of the time is spent on traversing the path-lists. This traversal is needed for inserting terms into, and deleting them from the index, as well as for computing intersection of path-lists in retrieval attempts. Computing intersection amounts to checking elements of the lists one by one. In any retrieval attempt, we are bound to traverse at least one of the lists completely. In the worst case, all of them are traversed completely. Similarly, when inserting a term identifier into a path list, we have to skip one by one all smaller elements of the list. When the length of the path lists grows with the growth of the indexed set, the cost of such traversal becomes inappropriately high.

In order to identify possible directions for optimisation, we consider the following example. For simplicity we assume that the term identifiers are integer numbers. Suppose that we want to intersect the list $[11]$ with the list $[1, \dots, 13]$, i.e., a list, containing all integers between 1 and 13. Before we come to the right position in the second list, we will have to check all its ten elements that are smaller than 11. This could probably be done faster if we were able to jump over several elements of the second list in one go. To this end we employ *skip lists*, see [Pugh 1990].

Figure 8 shows an example of a skip list for representing the list of integers between 1 and 13. A skip list can be viewed as an ordered linked list whose nodes are equipped with additional pointers. In addition to the field for storing the key, a node in an ordered linked list contains a pointer to the next node. A node of a skip list may contain a whole array *next* of pointers to some nodes with greater values of keys. The size of the array is called the *height* of the node. The height of the node N_4 in Figure 8 is 3. The first element of the array, $next[0]$, always points at the next node in the list. If a pointer $next[i + 1]$ is not null, then $next[i]$ is also not null, and if $next[i + 1] \neq next[i]$, then the node identified by $next[i]$ precedes the one identified by $next[i + 1]$ in the list. The height of a node, identified by a pointer $next[i + 1]$, is always greater

or equal to i . Moreover, it is the nearest node in the list with such height.

The pointers $next[i]$, where $i > 0$, are intended to facilitate far jumps when the list is traversed. The pointer $N_4.next[2]$ in Figure 8 allows us to jump over three nodes: N_5, N_6 and N_7 . Normally, the bigger i is, the farther from the current node we can jump by using the pointer $next[i]$.

In addition to nodes of various heights, every skip list has an array *header* of pointers to nodes. The size of the array should not be smaller than the maximal height of a node in the list. Every cell $header[i]$ contains the pointer to the first node in the list whose height is greater than i . Thus, $header[0]$ always points to the very first node in the list.

Skip list is a probabilistic datastructure in the following sense. For optimal performance of search, the heights of the nodes are chosen randomly, but according to a certain distribution. Ideally, $\frac{1}{2}$ of all the nodes should be of the minimal height 1. Another $\frac{1}{4}$ of the nodes should be of the height 2. In general, the fraction of nodes with a height i should be $\frac{1}{2^i}$. Our implementation approximates this requirement by following the distribution when new nodes are added to a list. When a node is deleted, no adjustment is made to the remaining nodes. Since the height of a node is completely independent of the value of the key, we can not bias deletion towards some particular heights. Thus, deletion of nodes spoils the distribution very rarely.

Suppose now that we are looking for the first element in the skip list from Figure 8, with a key greater than or equal to 10. First, we check the key in the node identified by the highest pointer in the header, i.e., $header[3]$. The key $8 < 10$, which means that we can immediately go to the node N_8 , thus skipping 7 nodes in one step. The highest pointer in the node N_8 is null, so we have to go one level down in $N_8.next$. Now $N_8.next[2] \rightarrow key = 12 > 10$, and we have to go one level down again. Finally, $N_8.next[1] \rightarrow key = 10$ and we can report success. Only minor adjustments are needed to implement insertion and removal on the base of the search algorithm, as well as fast traversal needed for computing intersection of skip lists.

The expected cost of search in a skip list is logarithmic w. r. t. the length of the list (see [Pugh 1990]). This suggests that an accurate implementation of insertion, deletion and traversal on the base of the search algorithm may radically improve performance of the whole instance retrieval procedure. This hypothesis is supported by our experimental data. We have compared two implementations that differ only on the representation of path-lists on the PEQ problems. The implementation using skip lists turns out to be **12.6 times faster** on average than the standard one using ordinary sorted lists, although the index requires **1.51 more memory**. The memory consumption overhead can be explained by the need to store one additional pointer on average, compared to the nodes of ordinary sorted lists.

5 Rearranging the Order of Path-List Intersection

When running our implementations of path indexing on benchmarks with large indexed sets, we noticed that the path-lists participating in the intersection vary in size very much. Typically, the average size of the lists may be big, but at the same time there may be a few very short path-lists, corresponding to the paths that are rare in the indexed set. Long path-lists typically have relatively many elements in common, while short path-lists tend to have few. This observation gives rise to another optimisation.

Suppose we want to run our intersection procedure on the following five lists:

$$\begin{aligned}
L_1 &= [1, 2, 3, \dots, 100] && \text{(all integers from 1 to 100)} \\
L_2 &= [1, 3, 5, \dots, 99, 100] && \text{(odd integers from 1 to 99, and 100)} \\
L_3 &= [2, 4, \dots, 100] && \text{(even integers from 2 to 100)} \\
L_4 &= [1, 100] \\
L_5 &= [2, 100]
\end{aligned}$$

It turns out that the behaviour of our intersection procedure in Section 2.4.3 is very sensitive to the order in which the lists are arranged. First, let us see how *intersect*(L_1, L_2, L_3, L_4, L_5) works. In the beginning of the outer loop, it finds a common element in L_1 and L_2 , and makes it the current candidate. In our example, all the 51 elements of L_2 are eligible, and most of them are discarded at the next step as absent from the list L_3 . By the end of the procedure execution, all elements of L_1, L_2 and L_3 will be examined in some way. Representing the lists as skip lists does not help us much, since their main advantage, the possibility of far jumps, can not be used in these settings.

We can facilitate more far jumps if we *rearrange the order of lists, considering shorter lists before longer ones*. Execution of *intersect*(L_4, L_5, L_3, L_2, L_1) quickly discovers that $L_4 \cap L_5$ consists of only one element 100, thus in the lists L_3, L_2 and L_1 we only need to make a few far jumps skipping all elements smaller than 100. Moreover, it often happens that the intersection of a few shortest path-lists is empty, and we can avoid processing the longer lists completely.

Experimental comparison on the PEQ problems shows that adding this heuristic to the implementation based on skip lists, makes it **1.6 times** faster on average. Since the overhead for sorting the path-lists by size is very small, the implementation with this heuristic never performs worse than the one without it.

6 Compiled Cleanup

Now, when the retrieval of candidates has become relatively fast, a significant part of the overall time is often spent on filtering out the wrong candidates and extracting the substitutions, which makes it a potential target for optimisations. On cleanup, we call the procedure *MatchCand*(q, t) (see Section 2.4.3) on many different candidate terms t for one fixed query term q . This suggests that we can use *specialisation with compilation*.

The idea of *specialisation of algorithms with compilation*, see, e.g. [Voronkov 1995, Riazanov and Voronkov 2000], is based on *partial evaluation*. Suppose that we need to execute some algorithm $alg(A, B)$ in a situation where a value of A is fixed for potentially many different values of B . In order to do this efficiently, we can try to find a specialisation of alg for every fixed A , i.e., such a algorithm alg_A , that executing $alg_A(B)$ is equivalent to executing $alg(A, B)$. The specialised algorithm may be more efficient than the generic one, since it can exploit some particular properties of the fixed value A . Typically, $alg_A(B)$ can avoid some operations that $alg(A, B)$ would have to perform, if they are known to be redundant for this particular parameter A . In particular, we can often identify some tests that are true or false for A , unroll loops and recursion, etc.

The specialised algorithm has to be represented in a form that can be interpreted, usually as a code of some abstract machine. Then the code itself can be additionally optimised by equivalent transformations that rely only on the semantics of instructions of the abstract machine.

Instructions of the abstract machine can usually be represented as records. One field of such a record stores an integer tag that identifies the instruction type, other fields may be used for

storing additional parameters of the instruction, for example a pointer to another instruction representing a label, if the semantics of the instruction requires a jump. All instructions of a code can be stored in an array, or list, or tree. Interpretation is done by fetching instructions in some order, identifying their type and executing the actions associated with this type. In C or C++ we can use a **switch** statement to associate some actions with different instruction tags. Modern compilers usually compile a **switch** statement for an integer value with many cases rather efficiently. This is achieved by storing the address of the statement corresponding to a value i in the i -th cell of a special array.

Now, let us show how this method works for candidate cleanup. Suppose we are trying to match a query term $q = f(g(g(a, a), a), h(x_0), g(x_1, x_0))$ against some candidate term t . If the matching is successful, t must contain all paths in q and be of the form $f(g(g(a, a), a), h(t^{(1)}), g(t^{(2)}, t^{(3)}))$. Execution of the straightforward algorithm $MatchCand(q, t)$ requires traversing the terms q and t in parallel, and q has to be traversed completely. In particular, we have to traverse the subterm $g(g(a, a), a)$ both in q and t , which is redundant because we already know that t contains all paths of q . During traversal, we examine the top symbol of every visited subterm to see if it is a variable. This is also redundant, since we know the structure of q . Finally, for every occurrence of a variable in q we have to check if it has been instantiated. When we get to the first occurrence of x_0 in q , this test is redundant since we know that so far x_0 has not been considered at all. When we get to the second occurrence, the test is redundant again, since we know that x_0 has been instantiated when its first occurrence was visited.

These redundant actions can be avoided in the specialised version of $MatchCand$. The essence of the algorithm for $MatchCand_q(t)$ can be described in the following way. The procedure has to get to the subterms $t^{(1)}$, $t^{(2)}$ and $t^{(3)}$ as quickly as possible. As soon as we get both $t^{(1)}$ and $t^{(3)}$, we check that they are equal. When either $t^{(1)}$ or $t^{(3)}$ is known, we make it the instance of x_0 in the substitution θ , and as soon as $t^{(2)}$ is known, we make it the instance of x_1 .

Let us see how $MatchCand_q(t)$ can be compiled. We rely on the following representation of indexed terms. A term $f(t_1, \dots, t_n)$ is represented by a structure, containing f as the first field, together with an array *arg* of pointers identifying the subterms t_i . We assume that the terms are perfectly shared, that is, only one copy of every term is stored, then comparison of terms can be implemented by comparing the pointers. The memory of our abstract machine will consist of a variable A for storing a pointer to an array of term arguments, and a sufficiently large array B for storing such pointers. The resulting substitution θ is stored as an array *subst*, such that $subst[i]$ is the pointer representing $x_i\theta$. Parameters of instructions denoted by p and r will be pointers to cells in B and *subst* correspondingly, and those denoted by i are integer numbers. The list of basic instructions, together with their semantics, is given in Figure 9.

Now, the algorithm for $MatchCand_{f(g(g(a,a),a),h(x_0),g(x_1,x_0))}(t)$ can be compiled into the sequence of instructions, shown in Figure 10. In its turn, the code can be optimised in several ways. For example, instructions 6 and 7 can be swapped. This transformation makes sense because instruction 7 may fail, and in this case we do not have to execute instruction 6. Also, we can pack some subsequent instructions into bigger ones. For example, we can pack *Store p* and *Down i* into an instruction *StoreAndDown p, i* with the following semantics: $*p := A$; $A := A[i] \rightarrow arg$. This enables our implementation language compiler to use a machine register for storing the value of A used in the assignment $*p := A$ and for accessing $A[i]$ later. Additionally, this transformation reduces the total number of instructions, and thus the time spent on moving between instructions and examining their tags.

To assess the proposed technique experimentally, we have implemented it on top of the

Tag	Parameters	Semantics
<i>Init</i>		$A := t \rightarrow arg$
<i>Store</i>	p	$*p := A$
<i>Restore</i>	p	$A := *p$
<i>Down</i>	i	$A := A[i] \rightarrow arg$
<i>Instantiate</i>	r, i	$*r := A[i]$
<i>Compare</i>	r, i	<u>if</u> $*r \neq A[i]$ <u>return</u> failure
<i>Success</i>		<u>return</u> success

Figure 9: Instructions for the candidate cleanup abstract machine

```

0  Init                                     /* assume  $t = f(g(g(a, a), a), h(t^{(1)}), g(t^{(2)}, t^{(3)}))$  */
1  Store       $B + 0$                        /*  $A := (g(g(a, a), a), h(t^{(1)}), g(t^{(2)}, t^{(3)}))$  */
2  Down       1                             /*  $B[0] := (g(g(a, a), a), h(t^{(1)}), g(t^{(2)}, t^{(3)}))$  */
3  Instantiate  $subst + 0, 0$                /*  $A := (t^{(1)})$  */
4  Restore      $B + 0$                        /*  $subst[0] := t^{(1)}$  */
5  Down       2                             /*  $A := (g(g(a, a), a), h(t^{(1)}), g(t^{(2)}, t^{(3)}))$  */
6  Instantiate  $subst + 1, 0$                /*  $A := (t^{(2)}, t^{(3)})$  */
7  Compare      $subst + 0, 1$                /*  $subst[1] := t^{(2)}$  */
8  Success                                     /* if  $t^{(1)} \neq t^{(3)}$  return failure */

```

Figure 10: Code for $MatchCand_{f(g(g(a,a),a),h(x_0),g(x_1,x_0))}(t)$

candidate retrieval implementation described in the previous section. The new implementation with compiled cleanup proved to be on average **1.2 times faster** than the older one with straightforward cleanup on the PEQ problems. The compilation-related overhead never leads to any noticeable slowdown.

7 Relational Path Indexing

So far we have only discussed some ways of efficiently implementing the same standard path indexing scheme, where retrieval of candidates containing all paths from the query is completely separated from the cleanup operations. Now we are going to describe a novel path indexing based technique, where the cleanup is tightly integrated with the candidate retrieval process. Roughly, the integration is done by interleaving path-list intersection operations with equality tests on subterms of the indexed terms. We would like to make it clear from the beginning that the main purpose of introducing the new technique is not only performance gains over the efficient implementation of standard path indexing, presented earlier. The proposed technique employs a very declarative description of retrieval with perfect filtering, and as such, it provides a logically clear and flexible base for more complex operations, like instance retrieval modulo commutativity and backward subsumption on multiliteral clauses, which will be discussed later. A detailed formal description of the proposed technique would be unnecessarily long, so we restrict ourselves to considering an example and exposing the logic behind it.

7.1 General Scheme for Retrieval

Every path π in a term t identifies a subterm t/π in t , whose top symbol coincides with the last symbol in the path. For example, $t/[f.1.g] = g(a, a)$ and $t/[f.3.h] = h(a, b)$, if $t = f(g(a, a), a, h(a, b))$. Denote by \arg_i , where i is a positive integer, the function on terms which returns the i th argument of a term. For example, $\arg_3(f(x_0, a, h(x_1))) = h(x_1)$. Suppose that we have a query $q = f(g(x_0, x_0), a, h(x_1, x_0))$. Instances of q and corresponding substitutions can be characterised by the following condition: $t = q\theta$ if and only if $[f.1.g], [f.2.a], [f.3.h] \in \text{Paths}(t)$ and $x_0\theta = \arg_1(t/[f.1.g]) = \arg_2(t/[f.1.g]) = \arg_2(t/[f.3.h])$, and $x_1\theta = \arg_1(t/[f.3.h])$. Thus, we can organise retrieval of instances of q from a term set \mathcal{I} by simply selecting all terms in \mathcal{I} satisfying this condition.

In order to give a finer description of the proposed technique, we need a notion of *path-relation* defined below, which replaces the notion of path-list. From now on, given a set of indexed terms \mathcal{I} and a path π , \mathcal{I}_π will denote a *path-relation* rather than a *path-list* for π . The word “relation” is used in the same sense as in relational databases, i.e., it means a finite collection of records with named attributes. If a path π ends in a function symbol f of arity n , the path-relation \mathcal{I}_π will have an attribute **term** ranging over terms (term identifiers), and attributes **arg_i**, whose values are also terms. For any term $t \in \mathcal{I}$, if the path π is in t and $t/\pi = f(t_1, \dots, t_n)$, then the path-relation will contain such a record r , that $r.\text{term} = t$ and $r.\text{arg}_1 = t_1, \dots, r.\text{arg}_n = t_n$.

As an example, consider a small indexed set $\mathcal{I} = \{t_1, t_2, t_3\}$, where $t_1 = f(g(a, a), a, h(b, a))$, $t_2 = f(g(a, b), a, h(b, b))$ and $t_3 = f(g(a, a), a, a)$. The path-relations for the paths $[f.1.g]$, $[f.2.a]$ and $[f.3.h]$ are shown in Figure 11.

The analogy with relational databases can be extended to retrieval in the following way. Since \mathcal{I} can be viewed as a relational database, we can describe the set of all substitutions corresponding to instances of a query q , as an SQL query. For example, the query $q =$

$\mathcal{I}_{[f.1.g]} :$	<table><tr><th>term</th><th>arg₁</th><th>arg₂</th></tr><tr><td>t_1</td><td>a</td><td>a</td></tr><tr><td>t_2</td><td>a</td><td>b</td></tr><tr><td>t_3</td><td>a</td><td>a</td></tr></table>	term	arg ₁	arg ₂	t_1	a	a	t_2	a	b	t_3	a	a
term	arg ₁	arg ₂											
t_1	a	a											
t_2	a	b											
t_3	a	a											
$\mathcal{I}_{[f.2.a]} :$	<table><tr><th>term</th></tr><tr><td>t_1</td></tr><tr><td>t_2</td></tr><tr><td>t_3</td></tr></table>	term	t_1	t_2	t_3								
term													
t_1													
t_2													
t_3													
$\mathcal{I}_{[f.3.h]} :$	<table><tr><th>term</th><th>arg₁</th><th>arg₂</th></tr><tr><td>t_1</td><td>b</td><td>a</td></tr><tr><td>t_2</td><td>b</td><td>b</td></tr></table>	term	arg ₁	arg ₂	t_1	b	a	t_2	b	b			
term	arg ₁	arg ₂											
t_1	b	a											
t_2	b	b											

Figure 11: Path relations for $\mathcal{I} = \{t_1, t_2, t_3\}$, where $t_1 = f(g(a, a), a, h(b, a))$, $t_2 = f(g(a, b), a, h(b, b))$ and $t_3 = f(g(a, a), a, a)$

select $\mathcal{I}_{[f.1.g]}.term$ **instance**, $\mathcal{I}_{[f.1.g]}.arg_1$ \mathbf{x}_0 , $\mathcal{I}_{[f.3.h]}.arg_1$ \mathbf{x}_1
from $\mathcal{I}_{[f.1.g]}$, $\mathcal{I}_{[f.2.a]}$, $\mathcal{I}_{[f.3.h]}$
where
 $\mathcal{I}_{[f.1.g]}.term = \mathcal{I}_{[f.2.a]}.term = \mathcal{I}_{[f.3.h]}.term$
and
 $\mathcal{I}_{[f.1.g]}.arg_1 = \mathcal{I}_{[f.1.g]}.arg_2 = \mathcal{I}_{[f.3.h]}.arg_2$

Figure 12: Relational query for $f(g(x_0, x_0), a, h(x_1, x_0))$

$f(g(x_0, x_0), a, h(x_1, x_0))$ translates into the SQL query depicted in Figure 12. To evaluate this query, we build the Cartesian product of the path-relations, specified in the **from**-clause, throw away those records that do not satisfy the **where**-clause and, finally, build the answer table with two attributes corresponding to the columns $\mathcal{I}_{[f.1.g]}.term$, $\mathcal{I}_{[f.1.g]}.arg_1$ and $\mathcal{I}_{[f.3.h]}.arg_1$ in the Cartesian product. The value of the **instance** attribute in the answer table contains a retrieved instance, and the values for \mathbf{x}_0 and \mathbf{x}_1 represent the corresponding substitution. Of course, this is only a declarative view. One would almost certainly use a more efficient process in practice. The **where**-clause in the query contains a number of equality conditions on the attributes of the input tables. The conditions on the **term** attribute correspond to candidate retrieval, since they ensure that extracted indexed terms have all the paths from the query. The conditions on **arg_i** implement the necessary cleanup, as they filter out the candidates that have different terms under the paths leading to equal variables in the query. In the relational database terminology, such constraints are called *join conditions*, and the answer table is called a *join* of the input relations. The proposed technique, which we call *relational path indexing*, provides perfect filtering for instance retrieval, provided that all the necessary path-relations participate in the join and all the necessary conditions are present in the **where**-clause.

7.2 Implementations

In all our implementations of relational path indexing, the indexing datastructure is very similar to that for the standard variant of path indexing, described earlier. We are using the same trie to access the path-relations. The path-relations themselves are represented by skip lists, whose nodes contain an indexed term as the key for ordering, and an additional pointer to the array of immediate subterms of the subterm, identified by the path. Using the indexed term as a key allows efficient joins on the **term** attribute. They are computed in the same way as intersection of path-lists represented by skip lists. When a record of a path-relation has been joined to others by the **term** attribute, the additional pointer can be used to access the values of **arg_i**.

7.2.1 Simple Implementation of Retrieval

As a starting point, we have implemented retrieval in a way similar to the one used in the standard technique. The join on **term** and the cleanup are separated. Roughly speaking, we first compute (in a lazy manner) the join on **term**, and for each record from the join we check the cleanup conditions.

7.2.2 Early Cleanup

The relational description of a query is declarative, it does not impose any constraints on how the joins are computed. In particular, the order in which they are computed does not affect the result. We can exploit this by doing some cleanup operations before some operations needed to build the join on **term**. Let us consider the query $q = f(g(x_0, x_0), a, h(x_1, x_0))$ from the example above. As soon as we know that some term t has the path $[f.1.g]$, we can check the condition $arg_1(t/[f.1.g]) = arg_2(t/[f.1.g])$. If this inexpensive test fails, we avoid execution of operations needed to check that $[f.2.a], [f.3.h] \in Paths(t)$. Similarly, when both $[f.1.g]$ and $[f.3.h]$ are known to be in t , we can immediately check $arg_1(t/[f.1.g]) = arg_2(t/[f.3.h])$. In general, we check a cleanup condition as soon as the participating records become available as a result of joining on the **term** attribute.

7.2.3 Combining Joins with Compiled Cleanup

The example considered above hides one inherent drawback of relational path indexing. In the query $q = f(g(x_0, x_0), a, h(x_1, x_0))$ all variable occurrences are arguments of subterms, identified by *maximal* paths. This is not the case in general, so we can no longer restrict ourselves to considering only maximal paths. Any path in a query leading to a subterm containing some variables as immediate subterms, contributes to the join. For example, in $q = g(x_0, g(x_0, g(x_1, a)))$, instead of simply taking all terms from $\mathcal{I}_{[g.2.g.2.g.2.a]}$ as candidates, we have to access corresponding records in $\mathcal{I}_{[g]}$, $\mathcal{I}_{[g.2.g]}$ and $\mathcal{I}_{[g.2.g.2.g]}$, in order to find possible instantiations of the query variables. Although dealing with such paths amounts to relatively efficient search in skip lists, it is too often beaten on performance by compiled cleanup. It is natural to try combining the advantages of compiled cleanup and early cleanup. We have implemented this idea in the following opportunistic way. Only those occurrences of variables in the query that are under maximal paths, give rise to join conditions that can be checked early. No additional path-relations need to be considered to check these conditions. If there are other variables, the join built in such a way may contain bad candidates filtered out by the additional compiled cleanup in the same way as it is done for standard path indexing.

7.3 Experiments

We have compared experimentally the three implementations, presented here against our best implementation of standard path indexing, which uses compiled cleanup, on the PEQ problems. All the new implementations use the same amount of memory, which is on average 27% more than the amount used by the fast standard path indexing. The first implementation (without early cleanup) is also 1.35 times slower. This is still an encouraging result. We can afford a moderate drop in performance, since we are mostly aiming at finding a flexible base for future extensions. The second implementation, employing early cleanup, is only 1.14 times slower. The third implementation, combining early and compiled cleanup, is already 1.11 times faster

than the standard technique with compiled cleanup. This figure shows the sheer value of the early cleanup optimisation.

7.4 Instance Retrieval Modulo Commutativity

Now we are going to show how relational path indexing can be used for instance retrieval modulo commutativity of some binary functions. For simplicity, we assume that the only commutative function symbols in the signature is $+$. *Syntactic equality* $=_C$ on terms *modulo commutativity* is defined in the following way:

$$s =_C t \Leftrightarrow \begin{cases} s = t, \\ \text{or } s = f(s_1, \dots, s_n), t = f(t_1, \dots, t_n), f \neq +, s_i =_C t_i \text{ for } i = 1, \dots, n, \\ \text{or } s = +(s_1, s_2), t = +(t_1, t_2), s_1 =_C t_1, s_2 =_C t_2, \\ \text{or } s = +(s_1, s_2), t = +(t_1, t_2), s_1 =_C t_2, s_2 =_C t_1 \end{cases}$$

Any term, equal modulo commutativity to a term t , will be called a *rotation* of t .

Term t is an *instance* of s *modulo commutativity*, if $t =_C s\theta$ for some substitution θ . Substitutions θ_1 and θ_2 are *equivalent modulo commutativity*, if $x\theta_1 =_C x\theta_2$ for any variable x . Now *instance retrieval modulo commutativity* is the following algorithmic problem. Suppose that we are given a term set \mathcal{I} , whose elements are pairwise distinct modulo commutativity. For any query term q , find all such substitutions θ , that $q\theta =_C t$ for some term $t \in \mathcal{I}$. Obviously, we can restrict ourselves by considering only substitutions that are pairwise inequivalent modulo commutativity.

One can try to solve this problem by adapting an indexing technique for ordinary instance retrieval. For example, a possible solution is to store all rotations of terms from \mathcal{I} in the index. In most cases this is not acceptable, as the index grows very quickly: every term with n commutative symbols may potentially have 2^n different rotations. Alternatively, we can keep the index as it is, and try all rotations of the query for retrieval, provided the retrieval algorithm is adjusted to accomodate syntactic equality modulo commutativity. In this case retrieval becomes costly, as the number of its invocations grows drastically. Fortunately, the idea of relational path indexing allows us to formulate an elegant solution which is free from such drawbacks and is likely to be efficient in practice.

The basic idea behind the proposed technique is as follows. The arguments of a commutative functor can permute freely, so we do not distinguish the argument numbers, following commutative symbols in the paths. This observation is made in the conclusion of [McCune 1992], and is also used in [Graf 1996] for defining a filter for instance retrieval modulo associativity and commutativity. As long as only commutativity is concerned, we go slightly further by providing a perfect filter.

A path π is called *normalised*, if every argument number immediately following a commutative symbol in π is equal to 1. With every path π we associate a normalised path $norm(\pi)$, which is obtained from π by replacing the corresponding argument numbers by 1. Additionally, we consider *denormalising* bit vector $denorm(\pi)$, such that $denorm(\pi)[i] = 0$ if and only if the i -th occurrence of a commutative symbol in π is followed by 1. For example, $norm([+.2.+.1.+.2.a]) = [+.1.+.1.+.1.a]$ and $denorm([+.2.+.1.+.2.a]) = 101$. It is obvious that any path π can be restored from $norm(\pi)$ and $denorm(\pi)$.

The indexing datastructure for a term set \mathcal{I} consists of path-relations \mathcal{I}_π for every normalised path π . Every path-relation \mathcal{I}_π has the attribute **term** ranging over terms, and an attribute


```

select  $\mathcal{R}_1.\text{term}$  instance,  $\mathcal{R}_1.\text{arg } x_0$ ,  $\mathcal{R}_2.\text{arg } x_1$ 
from  $\mathcal{I}_{[+]}$   $\mathcal{R}_1$ ,  $\mathcal{I}_{[+.1.+]}$   $\mathcal{R}_2$ ,  $\mathcal{I}_{[+.1.+1.f]}$   $\mathcal{R}_3$ ,  $\mathcal{I}_{[+.1.+1.f.2.+1.a]}$   $\mathcal{R}_4$ ,  $\mathcal{I}_{[+.1.+1.f.2.+]}$   $\mathcal{R}_5$ 
where
   $\mathcal{R}_1.\text{term} = \mathcal{R}_2.\text{term} = \mathcal{R}_3.\text{term} = \mathcal{R}_4.\text{term} = \mathcal{R}_5.\text{term}$ 
  and
   $\mathcal{R}_1.\text{denorm}|_1 \neq \mathcal{R}_2.\text{denorm}|_1$ 
  and
   $\mathcal{R}_2.\text{denorm}|_1 = \mathcal{R}_3.\text{denorm}|_1$ 
  and
   $\mathcal{R}_2.\text{denorm}|_2 \neq \mathcal{R}_3.\text{denorm}|_2$ 
  and
   $\mathcal{R}_3.\text{denorm}|_2 = \mathcal{R}_4.\text{denorm}|_2 = \mathcal{R}_5.\text{denorm}|_2$ 
  and
   $\mathcal{R}_4.\text{denorm}|_3 \neq \mathcal{R}_5.\text{denorm}|_3$ 
  and
   $\mathcal{R}_1.\text{arg} =_C \mathcal{R}_3.\text{arg}_1$  and  $\mathcal{R}_2.\text{arg} =_C \mathcal{R}_5.\text{arg}$ 

```

Figure 15: Relational query for $+(x_0, +(x_1, f(x_0, +(a, x_1))))$

values of the **term** attribute can be collected into lists. Checking the cleanup join conditions early is possible at least for the paths not containing commutative symbols. Moreover, it is likely that early checks of the conditions on the **denorm** vectors can also help. Comparing prefixes of the **denorm** vectors can be done cheaply by using bitwise shifts, if the vectors are properly represented. Compiled cleanup is directly applicable to the paths, not containing commutative symbols, and for others the code can use the values of **denorm** to direct traversal of a candidate term. Testing syntactic equality of terms modulo commutativity should not cause many problems. In the simplest case, the indexed terms and their subterms can be stored in a normalised form, unique for all terms, equal modulo commutativity. Alternatively, we can associate with every term a pointer to its normal form.

7.5 Backward Subsumption on Multiliteral Clauses

Now we are going to show how the relational path indexing scheme can be adapted to solving another hard decision problem: backward subsumption. The *subsumption relation* is defined on arbitrary first-order clauses, which can be viewed as multisets of literals, in the following way. A clause C *subsumes* a clause D , if there exists a substitution θ , such that $C\theta$ is a submultiset of D . For example, the clause $p(x, y) \vee q(y, z)$ subsumes the clause $p(a, b) \vee q(b, c) \vee r$, and the required substitution is $\{x \rightarrow a, y \rightarrow b, z \rightarrow c\}$. We can reduce the search space in a saturation-based prover by applying the so-called *backward subsumption*. When a new clause is added to the current clause multiset, we can discard those old clauses that are subsumed by the new one.

As a retrieval condition, backward subsumption amounts to finding all clauses in a clause multiset \mathcal{I} , subsumed by a given query clause Q . This problem is related to instance retrieval, since every literal in the query has an instance in a subsumed clause. So, it is natural to try to adapt an instance retrieval technique to the new problem. Here we describe a solution to the backward subsumption problem, based on relational path indexing, as implemented in

$\mathcal{I}_{[p.1.f]} :$				$\mathcal{I}_{[\neg p.1.f]} :$			
cl	arg ₁	arg ₂	lit	cl	arg ₁	arg ₂	lit
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
C	a	b	1	C	c	d	3
C	b	c	2	⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮				

Figure 16: Some path-relation records for $C = p(f(a, b)) \vee p(f(b, c)) \vee \neg p(f(c, d))$

our system Vampire ([Riazanov and Voronkov 2002]). Due to its importance and difficulty, backward subsumption is the subject of a separate forthcoming publication. The description below is superficial and only intended to illustrate the flexibility of our approach.

First of all, let us show how the index for a clause set \mathcal{I} is constructed. The indexed clauses are considered as sequences, rather than multisets. Only one sequence representation of a clause is used for indexing, and the choice of it is irrelevant. The notion of a path changes slightly. We are going to speak only about paths in literals, so the first element in a path is now always a *literal header*, i.e., a predicate with polarity, rather than a function symbol. As in the case of instance retrieval, the index provides a mapping of all paths π into corresponding path-relations \mathcal{I}_π . The path-relations have an attribute **cl** ranging over clause identifiers, a attributes **arg_i** as in instance retrieval, and an integer-valued attribute **lit**. Let C be a clause $L_1 \vee \dots \vee L_n$. If a literal L_j in C contains a path π , the path-relation \mathcal{I}_π should have a record r with the following values: $r.\mathbf{cl} = C$, $r.\mathbf{lit} = j$ and $r.\mathbf{arg}_i = \text{arg}_i(L_j/\pi)$. For example, consider the clause $C = p(f(a, b)) \vee p(f(b, c)) \vee \neg p(f(c, d))$. The entries in $\mathcal{I}_{[p.1.f]}$ and $\mathcal{I}_{[\neg p.1.f]}$ generated by this clause are shown in Figure 16.

With such an index we can perform retrieval of subsumed clauses by following our general scheme: by translating query clauses into relational queries. For example, consider a query clause $Q = p(f(x_0), g(x_1)) \vee p(f(x_1), g(x_2)) \vee \neg p(x_0, a)$, which translates into the SQL query in Figure 17. The translation is justified by the following considerations. Every clause subsumed by Q must have three different literals $p(f(t^{(1)}, g(t^{(2)}))$, $p(f(t^{(3)}, g(t^{(4)}))$ and $p(t^{(5)}, a)$, where $t^{(1)} = t^{(5)}$ and $t^{(2)} = t^{(3)}$. As such, it must contain the paths $[p.1.f]$, $[p.2.g]$, $[\neg p]$ and $[\neg p.2.a]$, moreover the first and second paths must occur at least twice. This requirement translates into the join condition on the **cl** attribute. Two of the literals above have both $[p.1.f]$ and $[p.2.g]$, and the remaining one has both $[\neg p]$ and $[\neg p.2.a]$. This gives rise to the second line in the **where**-condition. All the tree literals are different elements of the subsumed clause (although, may be syntactically equal), and this is reflected by the third line in the **where**-condition. The last line realises the requirements $t^{(1)} = t^{(5)}$ and $t^{(2)} = t^{(3)}$.

8 More Experiments

To give the reader the big picture, we start this section with a summary of results of the experiments discussed in different parts of the report. The table in Figure 18 shows the resource usage by different implementations, relative to our implementation of discrimination trees. To give the reader a feeling of complexity of the used benchmarks, we also mention the average time and memory usage per problem. The names of techniques are abbreviated in the following way. When the standard scheme is used, which separates candidate retrieval from cleanup, the

```

select  $\mathcal{R}_1.\text{cl}$  subsumed_clause
from  $\mathcal{I}_{[p.1.f]}$   $\mathcal{R}_1$ ,  $\mathcal{I}_{[p.2.g]}$   $\mathcal{R}_2$ ,  $\mathcal{I}_{[p.1.f]}$   $\mathcal{R}_3$ ,  $\mathcal{I}_{[p.2.g]}$   $\mathcal{R}_4$ ,  $\mathcal{I}_{[\neg p]}$   $\mathcal{R}_5$ ,  $\mathcal{I}_{[\neg p.2.a]}$   $\mathcal{R}_6$ 
where
   $\mathcal{R}_1.\text{cl} = \mathcal{R}_2.\text{cl} = \mathcal{R}_3.\text{cl} = \mathcal{R}_4.\text{cl} = \mathcal{R}_5.\text{cl} = \mathcal{R}_6.\text{cl}$ 
  and
   $\mathcal{R}_1.\text{lit} = \mathcal{R}_2.\text{lit}$  and  $\mathcal{R}_3.\text{lit} = \mathcal{R}_4.\text{lit}$  and  $\mathcal{R}_5.\text{lit} = \mathcal{R}_6.\text{lit}$ 
  and
   $\mathcal{R}_1.\text{lit} \neq \mathcal{R}_3.\text{lit}$  and  $\mathcal{R}_1.\text{lit} \neq \mathcal{R}_5.\text{lit}$  and  $\mathcal{R}_3.\text{lit} \neq \mathcal{R}_5.\text{lit}$ 
  and
   $\mathcal{R}_1.\text{arg}_1 = \mathcal{R}_5.\text{arg}_1$  and  $\mathcal{R}_2.\text{arg}_1 = \mathcal{R}_3.\text{arg}_1$ 

```

Figure 17: Relational query for $p(f(x_0), f(x_1)) \vee p(f(x_1), f(x_2)) \vee \neg p(x_0, a)$

impl	$\frac{\text{time}(\text{disc})}{\text{time}(\text{impl})}$	$\frac{\text{memory}(\text{disc})}{\text{memory}(\text{impl})}$	average time(impl), sec	average memory(impl), Mb
discrimination trees	1.00	1.00	91.3	3.1
std	0.31	0.95	291.5	3.3
std+sl	3.95	0.63	23.1	4.9
std+sl+srtp	6.33	0.63	14.4	4.9
std+sl+srtp+ccu	7.58	0.63	12.1	4.9
rel	5.60	0.50	16.3	6.2
rel+ecu	6.64	0.50	13.8	6.2
rel+ecu+ccu	8.43	0.50	10.8	6.2

Figure 18: Experimental results on PEQ

name starts with **std**. If an implementation relies on skip lists (Section 4) rather than sorted ones, **sl** is added to the name, **srtp** means that sorting of path-lists (Section 5) is used, and **ccu** stands for “compiled cleanup” (Section 6). The implementations, based on relational queries (Section 7) contain **rel** in their names. Please, notice that every such implementation employs skip lists and sorting of path-relations. Early cleanup (Section 7.2.2) is indicated by **ecu**. The used implementation of discrimination trees is called **disc**.

So far we have only experimented with pure equality problems. It is natural to ask whether the results are transferable to arbitrary problems with equality. The answer is “not quite”, as shown in Figure 19. The table indicates that the performance gains from all our optimisations, except the use of skip lists, are rather modest compared to the ones in the case of pure equality problems. There is a simple explanation to this effect. The introduction of skip lists is the only optimisation that affects the speed of updates, the others improve only retrieval, often at the expense of more complex updates. The PEQ problems are characterised by relatively high frequency of retrieval attempts. On average, only 2.5 update operations are performed per retrieval attempt, while for GEQ the corresponding value is 24.2 (see Figure 7). On GEQ the share of index updates in the overall time for **std+sl** and **rel** is 34% and 43% correspondingly, while on PEQ it is 4% and 6%, leaving much more space for retrieval-oriented optimisations.

impl	$\frac{\text{time}(\text{disc})}{\text{time}(\text{impl})}$	$\frac{\text{memory}(\text{disc})}{\text{memory}(\text{impl})}$	average time(impl), sec	average memory(impl), Mb
discrimination trees	1.00	1.00	23.0	3.4
std	0.07	0.98	309.8	3.5
std+sl	2.94	0.63	7.8	5.3
std+sl+srtp	3.89	0.63	5.9	5.3
std+sl+srtp+ccu	4.33	0.63	5.3	5.3
rel	3.57	0.52	6.4	6.5
rel+ecu	3.91	0.52	5.9	6.5
rel+ecu+ccu	4.48	0.52	5.1	6.5

Figure 19: Experimental results on GEQ

Finally, we would like to warn the reader against drawing from our results a conclusion that path indexing is always preferable to discrimination trees. We can easily imagine a situation when memory usage or index maintenance time is critical. Discrimination trees is a very universal technique, which means that the same index structure can support more than one retrieval relation, thus saving on memory and updates. Substitution trees [Graf 1995] or context trees [Ganzinger, Nieuwenhuis and Nivela 2001] could be another technique worth trying.

9 Related and Future Work

The path indexing method was introduced by M. Stickel in [Stickel 1989] and R. Ramesh e.al. [Ramesh et al. 1990]. In [McCune 1992] W. McCune compares experimentally an implementation of path indexing, corresponding closely to what we call standard implementation, with a number of variants of discrimination trees. The comparison is done for several retrieval relations, including instance retrieval. In the book of P. Graf [Graf 1996], a wide range of path indexing techniques is thoroughly described for most, if not all, interesting retrieval relations, together with a number of implementation tricks. In particular, an unsuccessful attempt to combine candidate retrieval with cleanup in one procedure, *extended path indexing*, is discussed. Also, a promising technique is formulated for dealing with retrieval modulo associativity and commutativity.

The work, described in this report, can be continued in the following directions. First of all, we are preparing a thorough experimental assesment for the backward subsumption implementation outlined in Section 7.5, and, hopefully, a publication will follow, describing the technique in more detail. The technique for instance retrieval modulo commutativity (Section 7.4) is to be implemented, and a COMPIT-style infrastructure for experiments is to be created. We also believe that efficient indexing techniques can be constructed for instance retrieval modulo associativity and commutativity along the lines of relational path indexing. Another promising direction is to try to improve the behaviour of the variation of the standard technique, which restricts the length of a path (see [McCune 1992]) and, thus, uses less memory. We believe that some improvement may be obtained by using compiled cleanup as in Section 6.

References

- DERSHOWITZ N. AND PLAISTED D. [2001], Rewriting, in A. Robinson and A. Voronkov, eds, ‘Handbook of Automated Reasoning’, Vol. I, Elsevier Science, chapter 9, pp. 533–608.
- GANZINGER H., NIEUWENHUIS R. AND NIVELA P. [2001], Context trees, in R. Goré, A. Leitsch and T. Nipkow, eds, ‘First International Joint Conference on Automated Reasoning, IJCAR 2001’, Vol. 2083 of *Lecture Notes in Artificial Intelligence*, Springer Verlag, pp. 242–256.
- GRAF P. [1995], Substitution tree indexing, in J. Hsiang, ed., ‘Proceedings of the 6th International Conference on Rewriting Techniques and Applications (RTA-95)’, Vol. 914 of *Lecture Notes in Computer Science*, Kaiserslautern, pp. 117–131.
- GRAF P. [1996], *Term Indexing*, Vol. 1053 of *Lecture Notes in Computer Science*, Springer Verlag.
- MCCUNE W. W. [1992], ‘Experiments with discrimination-tree indexing and path indexing for term retrieval’, *Journal of Automated Reasoning* **9**(2), 147–167.
- NIEUWENHUIS R., HILLENBRAND T., RIAZANOV A. AND VORONKOV A. [2001], On the evaluation of indexing techniques for theorem proving, in R. Goré, A. Leitsch and T. Nipkow, eds, ‘First International Joint Conference on Automated Reasoning, IJCAR 2001’, Vol. 2083 of *Lecture Notes in Artificial Intelligence*, Springer Verlag, pp. 257–271. see also the COMPIT homepage www.mpi-sb.mpg.de/~hillen/compit/.
- NIEUWENHUIS R. AND RUBIO A. [2001], Paramodulation-based theorem proving, in A. Robinson and A. Voronkov, eds, ‘Handbook of Automated Reasoning’, Vol. I, Elsevier Science, chapter 7, pp. 371–443.
- PUGH W. [1990], ‘Skip lists: A probabilistic alternative to balanced trees’, *Communications of the ACM* **33**(6), 668–676.
- RAMESH R., RAMAKRISHNAN I. AND WARREN D. [1990], Automata-driven indexing of Prolog clauses, in ‘Seventh Annual ACM Symposium on Principles of Programming Languages’, San Francisco, pp. 281–291.
- RIAZANOV A. AND VORONKOV A. [2000], Partially adaptive code trees, in M. Ojeda-Aciego, I. de Guzmán, G. Brewka and L. Pereira, eds, ‘Logics in Artificial Intelligence. European Workshop, JELIA 2000’, Vol. 1919 of *Lecture Notes in Artificial Intelligence*, Springer Verlag, Málaga, Spain, pp. 209–223.
- RIAZANOV A. AND VORONKOV A. [2002], ‘The design and implementation of Vampire’, *AI Communications* **15**(2-3), 91–110.
- SEKAR R., RAMAKRISHNAN I. AND VORONKOV A. [2001], Term indexing, in A. Robinson and A. Voronkov, eds, ‘Handbook of Automated Reasoning’, Vol. II, Elsevier Science, chapter 26, pp. 1853–1964.
- STICKEL M. [1989], The path indexing method for indexing terms, Technical Report 473, Artificial Intelligence Center, SRI International, Menlo Park, CA.
- SUTTNER C. B. AND SUTCLIFFE G. [1997], The TPTP problem library - v2.1.0, Technical Report JCU-CS-97/8, Department of Computer Science, James Cook University.
*<http://www.cs.jcu.edu.au/ftp/pub/techreports/97-8.ps.gz>
- VORONKOV A. [1995], ‘The anatomy of Vampire: Implementing bottom-up procedures with code trees’, *Journal of Automated Reasoning* **15**(2), 237–265.